arXiv:1412.3095v1 [cs.CC] 9 Dec 2014

Scheduling with two non-unit task lengths is NP-complete

Jan Elffers Mathijs de Weerdt

December 10, 2014

Abstract

We consider the non-preemptive task scheduling problem with release times and deadlines on a single machine parameterized by the set of task lengths the tasks can have. The identical task lengths case is known to be solvable in polynomial time. We prove that the problem with two task lengths is NP-complete, except for the case in which the short jobs have unit task length, which was already known to be efficiently solvable.

1 Introduction

The problem considered in this paper is the non-preemptive task scheduling problem with release times and deadlines. In the three-field notation, the problem is denoted $1|r_i|L_{\max}$. In this offline scheduling problem, there is a set of tasks, each having a release time, a deadline and a processing time, that need to be scheduled on a single machine without preemption. The goal is to schedule the tasks without preemption such that no task starts before its release time and no task completes much later than its deadline. Formally, the $L_{\max}$ optimization criterion asks to minimize the maximum lateness, that is, the maximum difference in time between a task's completion time and its deadline. The problem is NP-complete in theory by an easy reduction from bin packing (Pinedo, 2008), but branch and bound algorithms work well in practice, at least on certain distributions of randomly generated instances of up to 1000 jobs (Carlier, 1982). The decision problem asks whether a schedule without late jobs exists. The formal definition is as follows. We require that all release times, deadlines and task lengths are integers.

Definition 1 (Non-preemptive single machine scheduling with release times and deadlines)**.** *Given a set of tasks $J = \{([r_i, d_i], p_i) \mid i = 1, \ldots, n\}$, where $r_i, d_i \in \mathbb{Z}$ are the task's release time and deadline, together forming the task's* availability interval *$[r_i, d_i]$, and $p_i \in \mathbb{N}$ is the task's processing time, does there exist a* schedule*, that is, an assignment of starting times $t : \{1, \ldots, n\} \to \mathbb{R}$ to the tasks, such that $r_i \leq t(i) \leq d_i - p_i$ for all $i = 1, \ldots, n$, and the set of execution intervals $\{[t(i), t(i) + p_i) \mid i = 1, \ldots, n\}$ is pairwise disjoint?*

We study here a parameterized version of the problem, with the parameter the set of task lengths $P(J) = \{p_i \mid (r_i, d_i, p_i) \in J\}$. The case $p_i = 1$ (unit task lengths) is solved by the greedy Earliest Due Date (EDD) algorithm; the general case $p_i = p$ (identical task lengths) can also be solved in $O(n \log n)$ time (Garey et al., 1981) though this algorithm is much more complicated. The generalization to $P = \{1, p\}$ can be solved using a linear programming formulation (Sgall, 2012). This approach computes a sequence of starting times such that the length-p jobs can be assigned to these starting times, with additional constraints that guarantee sufficient idle time for the unit length jobs. These additional constraints simply lower bound the amount of idle time left by the long jobs schedule over a number of intervals, as if the unit length jobs were preemptive. This is sufficient because release times, deadlines and processing times are integers so by discretization starting times can always be integers. If the smaller task length is non-unit, this does not work anymore.

The complexity status of the general two-task-lengths problem has been noted as an open problem (Simons and Warmuth, 1989; Sgall, 2012). In this paper we prove NP-completeness of the problem for any fixed pair of non-unit task lengths. Formally, we have the following result.

Theorem 2 (NP-completeness result)**.** *For any two fixed non-unit integer task lengths $p > q > 1$, the non-preemptive single machine scheduling problem with release times and deadlines on the set of task lengths $\{p, q\}$ is NP-complete.*

Notation: Throughout the paper we call length-p jobs *long* jobs and length-q jobs *short* jobs.

2 Overview of the reduction

Our proof of NP-completeness is via an auxiliary scheduling problem $\mathrm{AUX}(p, q)$ that is defined for any two integer task lengths (including the case $q = 1$). We prove this problem to be polynomial-time reducible to the original problem and, in the case both task lengths are non-unit, to be NP-complete. Together, this implies NP-completeness of the two non-unit task lengths case of the original problem, proving our main result. Formally, our lemma's are the following.

Lemma 3 (Reducibility of the auxiliary problem to the original problem)**.** *For any two integer task lengths $p > q \geq 1$, the problem $\mathrm{AUX}(p, q)$ is polynomial-time reducible to the original scheduling problem on task lengths $\{p, q\}$.*

Lemma 4 (NP-completeness of the auxiliary problem for non-unit task lengths)**.** *For any two* non-unit *integer task lengths $p > q > 1$, the problem $\mathrm{AUX}(p, q)$ is NP-complete.*

Our auxiliary problem is the following extension of the original problem. Next to a set of tasks J with an ordinary availability interval $[r, d]$, we add two sets of "pending" tasks J_p, J_q with two deadlines, one for each task length, which have equal size. We assume the ordinary tasks to have non-negative release times and choose $t = 0$ as the common release time of all pending tasks. Each pending task has two deadlines (d', d), where $d' \leq d$. Let $N = |J_p| = |J_q|$ be the number of tasks per length. We demand the pairs of deadlines per task length to form an ordered sequence $d'[1] \leq d[1] \leq d'[2] \leq d[2] \leq \ldots \leq d'[N] \leq d[N]$, and the long tasks to be relatively urgent compared to the short tasks: $d_p[i] \leq d'_q[i]$ for all $i = 1, \ldots, N$. The problem is to find a feasible schedule of $J \cup J_p \cup J_q$ in which at least one of the two pending tasks with index i, that is, $J_p[i]$ or $J_q[i]$, finishes by its early deadline d'. The formal definition is as follows.

Definition 5 (Problem $\mathrm{AUX}(p, q)$)**.** *Let $p > q \geq 1$ be two integer task lengths. Given are a set of jobs J with task lengths from $\{p, q\}$ and non-negative release times, and two sequences of pairs of deadlines $(d'_p, d_p)[1 \ldots N]$ and $(d'_q, d_q)[1 \ldots N]$ of the same length N satisfying the following conditions:*

$$d'_p[1] \leq d_p[1] \leq d'_p[2] \leq d_p[2] \leq \ldots \leq d'_p[N] \leq d_p[N]$$
$$d'_q[1] \leq d_q[1] \leq d'_q[2] \leq d_q[2] \leq \ldots \leq d'_q[N] \leq d_q[N]$$
$$d_p[i] \leq d'_q[i] \quad \textit{for all } i = 1, \ldots, N$$

Define the sets of jobs $J_p[1 \ldots N]$ and $J_q[1 \ldots N]$ with $J_p[i] = ([0, d_p[i]], p)$ and $J_q[i] = ([0, d_q[i]], q)$ for $i = 1, \ldots, N$.

The problem is: does there exist a feasible schedule for $J \cup J_p \cup J_q$ such that for each $i = 1, \ldots, N$, $J_p[i]$ completes at time $\leq d'_p[i]$ or $J_q[i]$ at time $\leq d'_q[i]$?

In the next two sections, we prove two properties of this problem, Lemma 3 and Lemma 4, together implying Theorem 2.

3 Reducibility of $\mathrm{AUX}(p, q)$ to the original problem

Our reduction of the auxiliary problem to the original problem is as follows. We call the reduction the *stacked scheduling problem*. For each pending job, we create two jobs with ordinary availability intervals, which we call the *inner* and the *outer* job. The inner job has the early deadline d' and

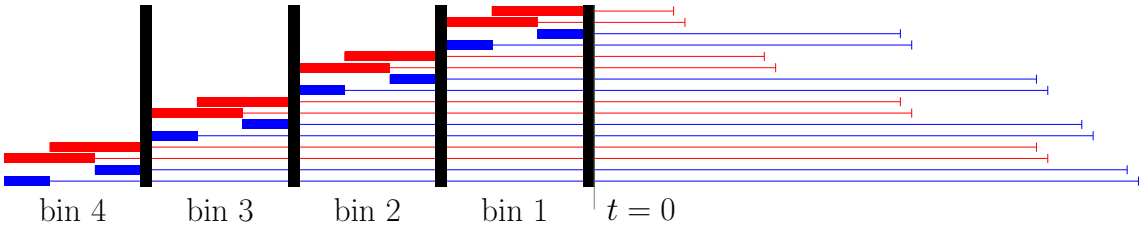

Figure 1: An instance of the stacked scheduling problem for $N = 4$ pending jobs per task length and without auxiliary jobs. The horizontal lines indicate availability intervals, and the horizontal bars indicate the earliest possible position at which the jobs can be scheduled. For the long (length-p) jobs they are colored red and for the short (length-q) jobs they are colored blue. The vertical bars denote separator jobs occupying fixed time intervals. Note that the availability intervals per task length form a nested sequence, and the long jobs have earlier deadlines than the short jobs for the same index.

the outer job has the late deadline d. The availability intervals of these jobs are extended to before $t = 0$ such that the intervals form a nested sequence per task length. We add *separator* jobs (jobs occupying fixed positions) before $t = 0$, partitioning the time line before $t = 0$ into *bins* of length $p + q$. We number the bins $1, \ldots, N$ starting from $t = 0$ backward in time. The release times of the pending jobs are aligned with these bins: the four jobs corresponding to index i have release times in the i'th bin; the outer jobs have release times starting at the earliest endpoint of the bin, and the inner jobs have release times such that they can only be scheduled at the end of their bin (or later). An example of the reduction is displayed in Figure 1. Formally, the reduction is defined as follows.

Definition 6 (Stacked scheduling problem)**.** *Let $p > q \geq 1$ be two integer task lengths. Given are a set of jobs J with task lengths from $\{p, q\}$ and non-negative release times, and two sequences of pairs of deadlines $(d'_p, d_p)[1 \ldots N]$ and $(d'_q, d_q)[1 \ldots N]$ of the same length N satisfying the conditions of Definition 5.*

Define $t_i = -(p + 2q) \cdot i + p + q$ for $i = 1, \ldots, N$. Let $J_{sep}[1 \ldots N]$ be the set of separator jobs, with $J_{sep}[i] = ([t_i, t_i + q], q)$ for $i = 1, \ldots, N$.

Define sets of jobs $J^I_p[1 \ldots N]$, $J^O_p[1 \ldots N]$, $J^I_q[1 \ldots N]$ and $J^O_q[1 \ldots N]$ as follows: for both task lengths $l = p, q$, $J^I_l[i] = ([t_i - l, d'_l[i]], l)$ and $J^O_l[i] = ([t_i - p - q, d_l[i]], l)$, for $i = 1, \ldots, N$.

The instance of the original problem is $J \cup J_{sep} \cup J^I_p \cup J^O_p \cup J^I_q \cup J^O_q$.

It is clear that the reduction works in polynomial time. We now prove correctness of the reduction. The idea is that a feasible schedule for the AUX(p, q) instance can be mapped to a feasible schedule of the reduced instance by copying the schedule of the ordinary jobs, placing inner/outer jobs at the positions where the corresponding jobs with two deadlines were scheduled, and placing the remaining inner/outer jobs before $t = 0$: we replace each pending job by the corresponding inner job, unless the completion time exceeds the early deadline; then we replace it by an outer job.

Lemma 7 (AUX(p, q) is feasible $\Rightarrow$ the reduced instance is feasible)**.** *Consider an instance of* AUX(p, q) *with parameters $(J, (d'_p, d_p)[1 \ldots N], (d'_q, d_q)[1 \ldots N])$. If the instance of* AUX(p, q) *has a feasible schedule, then the reduced instance also has a feasible schedule.*

Proof. The feasible schedule for the reduced instance is as follows. All jobs in J remain at the position specified by the schedule for the auxiliary problem; for the pending jobs, for both task lengths $l \in \{p, q\}$, if $J_l[i]$ completes by time $d'_l[i]$, place the i'th length-l inner job in its position; otherwise, place the i'th length-l outer job there. This way we have exactly one of the inner and outer job for each pending job scheduled after $t = 0$. We schedule the remaining inner/outer jobs before $t = 0$, placing the two remaining jobs with index i in bin i, for all $i = 1, \ldots, N$. Because in the original schedule, $J_p[i]$ or $J_q[i]$ finishes by its early deadline, and we place an inner job

after $t = 0$ for jobs finishing by their early deadline, there will be at least one outer job with index i to be scheduled in bin i. Therefore, we can place that job to the left, and the other job with index i to the right in the bin, which respects the release times. □

Lemma 8 (The reduced instance is feasible $\Rightarrow$ AUX(p, q) is feasible)**.** *Consider an instance of* AUX(p, q) *with parameters* $(J, (d'_p, d_p)[1 \dots N], (d'_q, d_q)[1 \dots N])$*. If the reduced instance has a feasible schedule, then the instance of* AUX(p, q) *also has a feasible schedule.*

Proof. We show an exchange argument that modifies any feasible schedule for the reduced instance such that for each index $i = 1, \dots, N$, bin i contains two jobs of index i, one short and one long job, with the job scheduled first being an outer job and the job scheduled last an inner job. This schedule implies a schedule for the AUX(p, q) instance, where the inner/outer jobs scheduled after $t = 0$ represent the pending jobs. Because each bin contains one outer job, for each index $i = 1, \dots, N$, one of $J_p[i]$ and $J_q[i]$ is represented by an inner job after $t = 0$, which means that this job finishes by its early deadline, as required.

Now, let a feasible schedule be given. Because all auxiliary jobs have release times after $t = 0$, they are not scheduled in the bins. We prove by induction that for $i = 1, \dots, N$, bins $1, \dots, i$ can be filled in this way. In step i, note that stack jobs with index i are scheduled in bin i or after $t = 0$, because by induction bins $j = 1, \dots, i - 1$ are filled with stack jobs with lower index. We say that we *swap in* a job into bin i if we add it to the left in the bin and push the other tasks in the bin to the right. This pushing to the right is possible because all stack task deadlines are after $t = 0$.

If bin i contains no long job, swap in the i'th outer long job which is scheduled after $t = 0$. Because the bin has length $p + q$, one short job can remain and the others, with total length at most p, fit in the original position of the long job. These jobs complete by their deadlines because $d'_q[i] \geq d_p[i]$ by definition, and the other short jobs have later deadlines. If bin i contains no short job, we can clearly swap in the i'th outer short job.

Now that bin i contains a long and a short job, we can swap the lefmost task in the bin to an i'th outer job, and the rightmost to an i'th inner job, because the availability intervals per stack are nested and those jobs are the ones with the smallest availability interval in their stacks that can be scheduled at these positions. □

4 NP-completeness of AUX(p, q) for $p > q > 1$

The second part of the proof is the NP-completeness result of the auxiliary problem for two non-unit task lengths. For readability, instead of defining the jobs in the reduction using only formulas, we define a model and express the reduction in terms of this model. This model consists of the sets of long and short pending jobs with two deadlines, and a set of auxiliary jobs with availability intervals of size $O(p)$, that is, constant for fixed p, q. The deadlines of the pending jobs are chosen such that they cannot start more than a constant ($O(p)$) time before their deadline in a feasible schedule. Therefore, we group jobs with almost the same deadlines. Our model is defined as a concatenation of *blocks* and separator jobs. A block is a set of auxiliary jobs and at most one pending job per task length. The release times and deadlines are defined relatively so that each block represents a local scheduling problem. Blocks are allocated an interval of length equal to the sum of task lengths of the tasks they define. We require the relative deadlines to be at most one time unit later than the length of each block so that we can force jobs from earlier blocks to be scheduled before jobs from later blocks.

The separators between blocks are placed one time unit after the end of the allocated time interval of the block immediately preceding them. Because both task lengths are non-unit, this unit of time cannot be filled by tasks from later blocks. The role of the separators is the following: if the last job from one block finishes one time unit after the block's allocated time interval, then each succeeding block F must be scheduled without idle time on $[1, L(F) + 1]$, where $L(F)$ is the block's length, until a separator is encountered. This allows one to do this propagation independently at multiple time intervals. We now formally define the model.

Definition 9 (Block)**.** *A block is defined as a tuple* $(J_{aux}, (d'_p, d_p), (d'_q, d_q))$*. Here* J_{aux} *is a set of auxiliary jobs with non-negative release times,* $d'_p \leq d_p$ *together represent a length-p pending job with two deadlines, and* $d'_q \leq d_q$ *together represent a length-q pending job with two deadlines. Both the long job and the short job with two deadlines are optional. The length L of the block is the sum of task lengths defined in it. All deadlines defined in the block must be in* $[1, L+1]$*.*

Definition 10 (Filled AUX(p, q) problem)**.** *Given are a sequence of blocks* $F = (F_1, \ldots, F_{|F|})$ *that defines N pending jobs per task length and a set of indices of blocks* I_{sep} *after which a separator is placed. Let* $s_{p,1} < s_{p,2} < \ldots < s_{p,N}$ *and* $s_{q,1} < s_{q,2} < \ldots < s_{q,N}$ *be the indices of blocks defining long and short stack jobs, respectively. We require that the i'th length-p job is defined before the i'th length-q job:* $s_{p,i} < s_{q,i}$ *for all* $i = 1, \ldots, N$*. Define offsets for the blocks as follows:*

$$\text{ofs}(F_i) = \begin{cases} 0 & i = 1 \\ \text{ofs}(F_{i-1}) + L(F_{i-1}) + q + 1 & \text{if } i - 1 \in I_{sep} \\ \text{ofs}(F_{i-1}) + L(F_{i-1}) & \text{otherwise} \end{cases}$$

Define separator jobs J_{sep} *and auxiliary jobs* J_{aux} *as follows:*

$$J_{sep} = \{([\text{ofs}(F_i) + L(F_i) + 1, \text{ofs}(F_i) + L(F_i) + q + 1], q) \mid i \in I_{sep}\}$$
$$J_{aux} = \bigcup_{i=1,\ldots,|F|} \{([r + \text{ofs}(F_i), d + \text{ofs}(F_i)], l) \mid (r, d, l) \in J_{aux}(F_i)\}$$

Define the deadlines $(d'_p, d_p)[1 \ldots N]$ *and* $(d'_q, d_q)[1 \ldots N]$ *of the pending jobs as follows: for both task lengths* $l \in \{p, q\}$*, for* $i = 1, \ldots, N$*,* $d'_l[i] = d'_l(F_{s_{l,i}}) + \text{ofs}(F_{s_{l,i}})$ *and* $d_l[i] = d_l(F_{s_{l,i}}) + \text{ofs}(F_{s_{l,i}})$*.*

The AUX(p, q) *instance is then given as* $(J_{sep} \cup J_{aux}, (d'_p, d_p)[1 \ldots N], (d'_q, d_q)[1 \ldots N])$*.*

This problem satisfies the conditions of Definition 5: because the relative deadlines in each block F_i are between 1 and $L(F_i) + 1$, deadlines of tasks in later blocks can never be earlier than deadlines of tasks in earlier blocks. Together with the requirement that the i'th long job is defined before the i'th short job, this means that the pending jobs satisfy the requirements.

We now state the claim that each job must be scheduled near its deadline formally.

Lemma 11. *Let* $((F_1, \ldots, F_{|F|}), I_{sep})$ *be an instance of the filled* AUX(p, q) *problem. For each block* F_i*, jobs defined in* F_i *cannot start before time* $\text{ofs}(F_i)$*.*

Proof. Let $s_1 < s_2 < \ldots < s_{|I_{sep}|}$ denote the set I_{sep}. We prove the statement by induction, first for blocks $1, \ldots, s_1$, then for $s_1 + 1 \ldots, s_2$, until $s_{|I_{sep}|} + 1, \ldots, |F|$. Note that for $i = 1, \ldots, s_1$, $\text{ofs}(F_i)$ equals $\sum_{j<i} L(F_j)$ and the latest deadline in blocks $F_1, \ldots, F_{i-1}$ is at most $\text{ofs}(F_i) + 1$. So tasks from F_i can overlap with $[0, \text{ofs}(F_i) + 1]$ for at most one time unit. Because $p, q > 1$, tasks from F_i have to start after $\text{ofs}(F_i)$. The first separator block starts at time $T = 1 + \sum_{j \leq s_1} L(F_j)$, so there is exactly one unit of idle time in $[0, T]$ left over by the first s_1 blocks, which cannot be filled because $p, q > 1$. Now the offset of the next block equals the end of the separator job's time interval, so we can repeat this argument. □

4.1 The reduction

Now we are ready to state the concrete reduction. We reduce from the boolean satisfiability (SAT) problem which is well-known to be NP-complete. Assume we have a SAT instance on n variables $x_1, x_2, \ldots, x_n$ with m clauses $C_1, C_2, \ldots, C_m$. We construct a filled AUX(p, q) problem as in Definition 10 with $N = n + m \cdot (2n + 1)$ pending jobs per task length. We label the pending jobs with v_i for the (only) pending job per variable x_i, and $c_{j,k}$ for the k'th pending job per clause C_j $(k = 0, \ldots, 2n)$. The order is:

$$c_{1,0}, \ldots, c_{m,0},$$
$$v_1, c_{1,1}, \ldots, c_{m,1}, c_{1,2}, \ldots, c_{m,2},$$
$$v_2, c_{1,3}, \ldots, c_{m,3}, c_{1,4}, \ldots, c_{m,4},$$
$$v_n, c_{1,2n-1}, \ldots, c_{m,2n-1}, c_{1,2n}, \ldots, c_{m,2n}$$

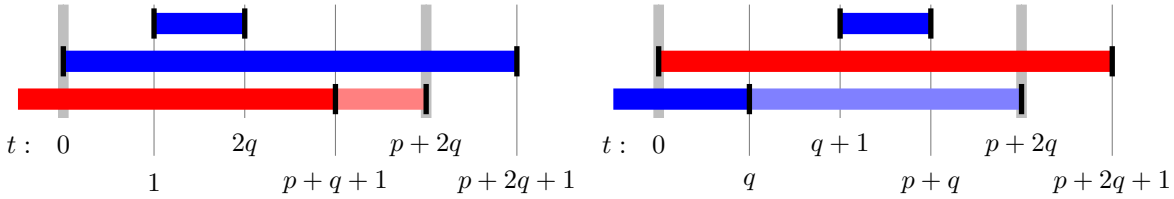

Figure 2: The two literal blocks: left, V^+ for positive literals and right, V^- for negative literals. The thick vertical lines are the start and end offset, together defining the length. The bars represent availability intervals: the red bars are for long jobs, the blue bars for short jobs. The jobs with release time before $t = 0$ and two deadlines denote pending jobs. Note that in both cases, if one demands the pending job to be completed by its early deadline, then the job scheduled last must complete at time $p + 2q + 1$, while completion time $p + 2q$ is possible if the pending job must complete only by its late deadline.

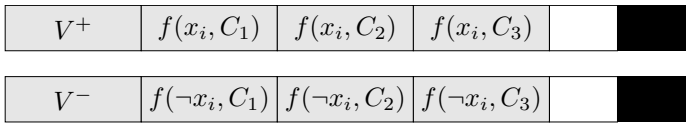

Figure 3: The main part of the reduction. Displayed are the sequences of blocks and a separator for x_i (top) and $\neg x_i$ (bottom) for one variable. In this example, there are three clauses. The gray boxes indicate the allocated time intervals for the blocks; function $f(l, C_j)$ gives a C_{active} block if $l \in C_j$ and a $C_{inactive}$ block otherwise. The black box is a separator job; the white box in front of it is a unit length empty space in the allocation.

The four non-trivial blocks we use are denoted V^+, V^-, C_{active} and $C_{inactive}$. V^+ and V^- represent positive and negative literals, respectively. V^+ defines $(d'_p, d_p) = (p + q + 1, p + 2q)$ and has auxiliary jobs $\{([1, 2q], q), ([0, p + 2q + 1], q)\}$; V^- defines $(d'_q, d_q) = (q, p + 2q)$ and has auxiliary jobs $\{([q + 1, p + q], q), ([0, p + 2q + 1], p)\}$. The two blocks are visualized in Figure 2. C_{active} and $C_{inactive}$ represent clauses and define a long and a short pending job and no auxiliary jobs. In both blocks, the short and long jobs have the same deadlines. For the C_{active} block, $(d'_p, d_p) = (d'_q, d_q) = (p + q, p + q + 1)$ and for $C_{inactive}$, $(d'_p, d_p) = (d'_q, d_q) = (p + q - 1, p + q + 1)$.

The sequence of blocks is the following. First, there are m "dummy" blocks without auxiliary jobs and only a long pending job $(d'_p, d_p) = (p - 1, p)$. Then there are $2n$ sequences of blocks for literals $l = x_1, \neg x_1, \ldots, x_n, \neg x_n$, in this order, each consisting of a "literal" block V^+/V^-, then m "clause" blocks $C_{active}/C_{inactive}$, and ended by a separator. The literal block is a V^+ if l is positive and a V^- if l is negative. The j'th clause block is C_{active} if $l \in C_j$ and $C_{inactive}$ otherwise. Finally, there are another m "dummy" blocks without auxiliary jobs and only a short pending job $(d'_q, d_q) = (q - 1, q)$.

The reader can verify that if the blocks are concatenated in this order, then the correspondence between pending job labels and blocks is the following: the first m long jobs $c_{1,0}, \ldots, c_{m,0}$ are defined in the first m dummy blocks and the last m short jobs $c_{1,2n}, \ldots, c_{m,2n}$ are defined in the last m dummy blocks; in the sequence of blocks for literal l, if l is positive, write $l = x_i$, then the V^+ block defines long job v_i and otherwise the V^- block defines short job v_i. The m clause blocks following the l'th literal block define long pending jobs $c_{1,l}, \ldots, c_{m,l}$ and short pending jobs $c_{1,l-1}, \ldots, c_{m,l-1}$.

The construction per variable is shown in Figure 3. The only blocks containing auxiliary jobs are the V^+ and V^- blocks. These blocks come in pairs: for each variable x_i, the v_i'th long pending job is defined in a V^+ block at the head of a sequence of clause blocks for literal x_i, and the v_i'th short pending job is defined in a V^- block at the head of the next sequence of clause blocks, for literal $\neg x_i$. If one wants the pending job to end by its early deadline in a V^+ or V^- block, then the last job in the block completes after the allocated time interval (see Figure 2). Because the V^+

and V^- blocks are linked per variable, one of the pending jobs must use the early deadline so in one of the sequences of clause blocks, all blocks are scheduled starting one unit after their offset, inducing constraints on the pending jobs defined in these clause blocks. The choice whether to let the long or the short v_i'th pending job finish by its early deadline corresponds to setting x_i to FALSE or TRUE, respectively.

It is obvious that this can be constructed in polynomial time. The resulting instance has $O(nm)$ jobs which is polynomial in the size of the SAT formula. We now prove the correctness of the reduction.

Lemma 12 (SAT model ⇒ feasible schedule)**.** *Let $v(x_1), v(x_2), \ldots, v(x_n)$ be a model for the SAT formula. Then the constructed* AUX(p, q) *instance has a feasible schedule.*

Proof. Schedule the first m and the last m pending jobs to start at the offsets of their blocks. For each variable x_i, $i = 1, \ldots, n$, schedule the V^+ and V^- blocks as follows. If $v(x_i) =$ TRUE, schedule the v_i'th long pending job at time $2q$ relative to its offset and the v_i'th short pending job at time 0 relative to its offset. The long pending job ends by its late deadline and the short pending job ends by its early deadline, so the V^+ block can be scheduled to complete at relative time $p + 2q = L(V^+)$, while the V^- block must complete at relative time $p + 2q + 1 = L(V^-) + 1$. If $v(x_i) =$ FALSE, schedule the v_i'th long pending job at time $q + 1$ relative to its offset and the v_i'th short pending job at time $p + q$ relative to its offset. The long pending job now ends by its early deadline and the short pending job ends by its late deadline, so the V^+ block must now complete at relative time $p + 2q + 1 = L(V^+) + 1$, while the V^- block can complete at relative time $p + 2q = L(V^-)$.

Schedule the clause blocks greedily as early as possible so that jobs from clause blocks start to be scheduled at their block's offset, or one time unit later if the corresponding V^+/V^- block was scheduled with one unit delay. For each clause C_j, $j = 1, \ldots, m$, let l_j be the index of a literal that satisfies the clause under v: positive literals x_i have index $l_j = 2i - 1$ and negative literals $\neg x_i$ have index $l_j = 2i$. Schedule the $2n$ clause blocks for C_j as follows: schedule the short job first at indices $1, \ldots, l_j - 1$, and the long job first at indices $l_j + 1, \ldots, 2n$. The order at index l_j does not matter: because there is no delay (as l_j is assigned TRUE by definition), both jobs finish by their early deadline $p + q$ in the C_{active} block. In the other indices, at least the job scheduled first completes by its early deadline, regardless of delay, because the early deadlines are at least $p + q - 1$ in both C_{active} and $C_{inactive}$. Because the late deadlines in the clause blocks are $p + q + 1$, the pending jobs will finish by their late deadlines anyway; it remains to be shown that for all $c_{j,k}$'th pairs of pending jobs, at least one job completes by its early deadline. For indices $1, \ldots, l_j$, the short job completes by its early deadline, satisfying the requirement for $k = 0, \ldots, l_j - 1$; for indices $l_j, \ldots, 2n$, the long job completes by its early deadline, satisfying the requirement for $k = l_j, \ldots, 2n$. Together, this means that for all pending job pairs at least one job finishes by its early deadline. □

Lemma 13 (Feasible schedule ⇒ SAT model)**.** *Suppose we have a feasible schedule for the constructed* AUX(p, q) *instance. Then there exists a model for the SAT formula.*

Proof. By Lemma 11, jobs from each block F_i cannot be scheduled before their block's offset ofs(F_i). In particular this means that in a V^+/V^- block, if one demands the pending job to complete by its early deadline, then the last job in the block completes at relative time $p + 2q + 1$, introducing a unit of idle time (see Figure 2). Furthermore, for each literal l, the jobs from the m corresponding clause blocks are scheduled in order after the V^+/V^- block corresponding to l. We choose a model v as follows: for each variable x_i, $i = 1, \ldots, n$, if the v_i'th long pending job completes after its early deadline, set $v(x_i) =$ TRUE; if the v_i'th short pending job completes after its early deadline, set $v(x_i) =$ FALSE (if both complete by their early deadline, set $v(x_i)$ arbitrarily).

We now argue that v satisfies all clauses. Suppose for the sake of contradiction that a clause C_j is falsified by v. For literals $l \notin C_j$, C_j is represented by a $C_{inactive}$ block which has early deadlines $p + q - 1$ so each block can only let one pending job finish by its early deadline. For literals $l \in C_j$, C_j is represented by a C_{active} block with early deadlines $p+q$; because $v(l) =$ FALSE,

the pending job corresponding to literal l must end by its early deadline, so it introduces a unit of delay in its V^+/V^- block, so again each block can only let one pending job finish by its early deadline. Finally the j'th block in the m dummy blocks at the start and the end of the schedule cannot let the pending job complete by its early deadline. So in total, there are $2n+1$ pairs of pending jobs with indices $c_{j,0}, c_{j,1}, \ldots, c_{j,2n}$, and only $2n$ of these pending jobs can finish by their early deadline, which gives a contradiction. □

5 Conclusion

We proved NP-completeness of the decision version of $1|r_i|L_{\max}$ with integer release times and deadlines if the set of task lengths is restricted to two non-unit integer task lengths, thereby proving that the case $\{1, p\}$ is the maximally theoretically solvable case. Another way to think about the $\{1, p\}$ problem is that the unit length jobs are actually preemptive jobs: the linear programming formulation of Sgall (2012) only works because the long jobs never need to preempt the unit length jobs. Therefore, we can consider a generalized problem with both preemptive and non-preemptive tasks and arbitrary fractional release times and deadlines. In terms of this problem, our result can be stated as follows: with respect to the set of task lengths, the only solvable case is the one in which *the non-preemptive tasks all have the same length.*

6 Acknowledgements

The authors thank Sicco Verwer for useful discussions.

References

Jacques Carlier. The one-machine sequencing problem. *European Journal of Operational Research*, 11(1):42 – 47, 1982. ISSN 0377-2217.

M. R. Garey, David S. Johnson, Barbara B. Simons, and Robert Endre Tarjan. Scheduling unit-time tasks with arbitrary release times and deadlines. *SIAM J. Comput.*, 10(2):256–269, 1981.

Michael L. Pinedo. *Scheduling: Theory, Algorithms, and Systems*. Springer Publishing Company, Incorporated, 3rd edition, 2008. ISBN 0387789340, 9780387789347.

Jiri Sgall. Open problems in throughput scheduling. In *ESA*, pages 2–11, 2012.

Barbara B. Simons and Manfred K. Warmuth. A fast algorithm for multiprocessor scheduling of unit-length jobs. *SIAM J. Comput.*, 18(4):690–710, 1989.